



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Warwick Economics Research Paper Series

How Transparency Kills Information Aggregation: theory and Experiment

Sebastian Fehrler & Niall Hughes

December, 2015

Series Number: 1088

ISSN 2059-4283 (online)

ISSN 0083-7350 (print)

This paper also appears as [CRETA Discussion Paper No: 2](#)

How Transparency Kills Information Aggregation: Theory and Experiment*

Sebastian Fehrler

University of Konstanz, and IZA

sebastian.fehrler@uni.kn

Niall Hughes

University of Warwick

N.E.Hughes@warwick.ac.uk

September 2015

Abstract

We investigate the potential of transparency to influence committee decision-making. We present a model in which career concerned committee members receive private information of different type-dependent accuracy, deliberate and vote. We study three levels of transparency under which career concerns are predicted to affect behavior differently, and test the model's key predictions in a laboratory experiment. The model's predictions are largely borne out - transparency negatively affects information aggregation at the deliberation and voting stages, leading to sharply different committee error rates than under secrecy. This occurs despite subjects revealing more information under transparency than theory predicts.

Keywords: Committee Decision-Making, Deliberation, Transparency, Career Concerns, Information Aggregation, Experiments, Voting, Strategic Communication.

JEL Classification Numbers: C92, D71, D83.

*We are especially grateful to Alessandra Casella for her support and many great discussions. We also wish to thank Peter Buisseret, Aniol Llorente-Saguer, Katja Michaelowa, Massimo Morelli, Becky Morton, Maik Schneider, Francesco Squintani, Jean-Robert Tyran and participants at several conferences and seminars for their valuable comments. This project was financially supported by the Swiss National Science Foundation (grant 100017_150260/1). All remaining errors are of course our own.

“In 1977, it was an article of faith in central banking that secrecy about monetary policy decisions was the best policy.” (Janet Yellen, Chair of the Federal Reserve)

“An institution that performs functions as vital as the Federal Reserve must operate in the public eye as much as possible.” (John C. Williams, President of San Francisco Federal Reserve, quote from 2011)

1. Introduction

Transparency in decision-making is a recurring and controversial topic in public debate. A recent example comes from the world of sport, where FIFA’s decision to hand the 2022 World Cup to Qatar - amid allegations of bribing - spurred calls for more transparency from soccer fans worldwide.¹ Another prominent debate is whether monetary policy committees such as the Federal Open Market Committee (FOMC) should be secretive or transparent (e.g., Williams (2012)). The quotes above show a move towards transparency over time, yet recent research has argued that too much transparency could be harmful (Meade and Stasavage (2008); Swank et al. (2008); Swank and Visser (2013)). This debate over transparency is not only a recent phenomenon. In the early 19th century Jeremy Bentham ([1816] 1999) argued for more transparency in parliamentary decision-making, while reservations had long been voiced by influential thinkers such as Thomas Hobbes and John Stuart Mill (see Stasavage (2007)).

The supposed boon of transparency is accountability; by making the decision-making process more transparent, this should align the incentives of the agents with those of the principals.² The downside of transparency is that if agents care about their reputations, they may pander to the principal, choosing an action which makes them appear smart but which is not necessarily in the principal’s interest (Prat (2005); Fox and Van Weelden

¹See, e.g., <http://www.theguardian.com/football/2014/sep/24/michael-garcia-fifa-investigation-public-world-cup>.

²A positive case for transparency is also made in the deliberative democracy literature (e.g., Habermas (1996) and Cohen (1996)) whose proponents view public deliberation as an essential element of legitimate democratic decision-making. See Landa and Meirowitz (2009) for a critical discussion of this literature.

(2012)). Indeed, recent theoretical literature has shown that reputational concerns influence committee behaviour differently under secrecy and transparency in the case of political committees (e.g. Stasavage (2007)), corporate boards (Malenko (2013)), and monetary policy committees (Levy (2007); Visser and Swank (2007); Swank and Visser (2013); Swank et al. (2008); Meade and Stasavage (2008); Gersbach and Hahn (2008)).

In the studies that allow for deliberation - an obviously realistic assumption for most committees - strategic communication is the key driver of these effects. However, recent experimental literature shows for a variety of settings that subjects are too truthful and too trusting as compared to theoretical 'cheap talk' predictions (Cai and Wang (2006); Wang et al. (2010); Battaglini and Makarov (2014)) and that a substantial fraction of subjects display psychological costs of lying (Gneezy (2005); Gneezy et al. (2013); Fischbacher and Föllmi-Heusi (2013)). In a study on committee voting, Goeree and Yariv (2011) even find that open deliberation leads to perfect information sharing in a set-up where committee members have private values regarding the committee decision and thus a clear incentive to keep their information private. These findings cast some doubt on the predictions from the theoretical literature on the role of career concerns in committees.

Unfortunately, empirical evidence on the effects of transparency on career concerned committee members is scarce and does not identify the mechanisms through which behavior is affected. Both Meade and Stasavage (2008) and Swank et al. (2008) present evidence of changes in deliberation in the FOMC after the decision to publish minutes of their meetings, but they present the same data as supporting evidence for their different career concerns models.³ The difficulty of identifying mechanisms is not surprising given the shortcomings of using field data. Firstly, it is almost certainly impossible to observe the level of ability, prior information and biases which committee members may have. To the extent that transparency may be interacting with these unobserved variables, it is impossible to identify the mechanism through which transparency affects behavior. Secondly, a controlled comparison of institutions is always very difficult as their different elements, e.g. the committee's voting rule, deliberation protocol and other factors are

³Another study is presented by Cross (2013), who does not test a particular model but compares behavior of members of the Council of the European Union under different levels of transparency and shows evidence suggesting that more extreme positions are taken under higher levels of transparency.

endogenous. For these reasons, we believe a laboratory experiment provides the cleanest way of examining the effects of transparency. Ours is the first laboratory experiment of committee decision-making with career concerns. By carefully controlling the information players have and varying the level of transparency, effects and mechanisms can be clearly identified.

We construct and test a model in which career concerns play out very differently under three levels of transparency. This set-up allows us to study how subjects react to the changing opportunities to act smart both at the deliberation and the voting stages of the decision process. As in most career concerns models, there are two types of committee members: high and low ability. First, committee members receive either a fully informative or a noisy signal about the true state of the world, depending on their ability. Then they have the opportunity to deliberate, and finally vote for or against changing the status quo in favor of an alternative option. A change of the status quo requires a unanimous vote of the committee. The first deliberation stage consists of a non-binding straw poll in which subjects can exchange information about their signals. This is followed by a second stage in which they can exchange information about their type, i.e. the signal strength. A committee member's utility depends on the principal's belief that he is of high ability. We study three different transparency regimes - *secrecy*, where votes and communication are secret and the principal only learns the committee decision, *transparency*, where both communication and individual votes are public, and the intermediate case of *mild transparency*, where individual votes are made public but deliberation is secret. Mild transparency thus reflects a committee practice of publishing individual voting records but not transcripts of the deliberation. This corresponds to the practice of the FOMC before 1993 (see Meade (2005)).⁴ We implement the three transparency regimes in a laboratory experiment in which the principal's belief about an agent's type is elicited with a proper scoring rule. This stated belief directly affects the agent's payoff. We implement the second deliberation stage as an open chat. We do this because it is more realistic if

⁴Mild Transparency also reflects the main structure of a scenario in which deliberation is nominally transparent but committee members arrange secretive pre-meetings for deliberation and then stage a public show meeting without real deliberation and public votes, which arguably describes what happened with decision-making in the FOMC after 1993 (Swank et al. (2008); Swank and Visser (2013)).

we want to approximate real committee meetings, and because free-form communication has been shown to matter for the level of strategic behavior in Goeree and Yariv (2011).

The following key insights from the model are corroborated in the experimental data. Firstly, committee members truthfully share all their information with each other under secrecy, but fail to do so under transparency. When the principal is watching, nobody wants to admit to being a low type, so information aggregation is incomplete. Secondly, incorrect group decisions are more prevalent under transparency than secrecy. This occurs for two reasons: (a) members are less truthful under transparency and thus aggregate poorer information, and (b) low ability members have an incentive to vote according to their private information even if they believe it to be wrong. This effect stems from the fact that if members change their position between the deliberation and voting stages, the principal can infer that they are a low ability type. Thirdly, a principal who strongly favours the status quo can actually be better off under transparency than secrecy. The failure of committee members to share information under transparency means they find it difficult to vote unanimously against the status quo. This means more mistakes when they should change the status quo, but less mistakes when they shouldn't. This asymmetry of errors appeals to a principal who is very concerned about wrongly changing the status quo. Here, transparency aids the principal *because* it hinders information aggregation in the committee. If committee members could privately communicate, as they can under secrecy, then a committee with a balanced prior (i.e. two low types with different signals) would implement each decision with 50% probability. Under transparency, without this ability, voting to signal ensures the status quo is always implemented. Finally, while the most informative equilibrium under mild transparency predicts behavior as under secrecy, the experiment shows aggregate level error rates quite similar to the transparency case. However, these result from quite different individual level behavior. Mild transparency leads to a significant level of deception at the first deliberation stage, whereas under transparency information aggregation breaks down at the second deliberation stage and at the voting stage. Under mild transparency, there are equilibria in which high types do not truthfully reveal their signals. This makes them easier to separate from low types in the

voting stage, resulting in a more favorable belief of the principal regarding their type.⁵ Our results suggest that the existence of these equilibria leads a number of high types to deliberate non-truthfully and thus to very different aggregate outcomes than predicted by the most informative equilibrium.

We observe a number of deviations from our theoretical predictions regarding individual level behavior. Principals are on average too optimistic in their assessment of committee members. However, they update their beliefs in the correct direction to all pieces of information they receive. Most committee members play a best response to the stated beliefs of the principals under secrecy and mild transparency, while many of them fail to best respond under transparency, where belief updating is more complicated and therefore also more complicated for committee members to anticipate. Also under transparency, open chat deliberation is more truthful and informative than predicted. This suggests that some subjects do indeed face psychological costs of lying, and confirms results from previous experimental studies. However, despite this truth-telling preference, the level of truth-telling varies greatly between treatments. Thus Goeree and Yariv's (2011) result on the power of open chat to moderate institutional differences does not extend to the case of transparency versus secrecy.

In the next section, we present and solve the model. We proceed with the experimental design and theoretical predictions for the chosen parameter values before discussing the aggregate and individual level results. We conclude with a discussion of the main results and their implications for the literature on career concerns in committees and cheap talk games.

2. The Model

A committee of two members must make a decision $D \in \{B(lue), R(ed)\}$ on behalf of a principal.⁶ The voting rule is unanimity: option R is implemented only if both members

⁵A structurally similar situation is described in Sobel (2013), section 2, in which he calls for experiments to study the predictive power of the most informative equilibrium in cheap talk games with additional equilibria which are preferred by at least one type of player.

⁶We present a two member model as we run two member games in the lab. In an n member committee players' behaviour would not change, though the probability of an incorrect group decision would.

vote for it, otherwise the status quo B is upheld. There are two equally likely states of the world $S \in \{B, R\}$ and each member i gets an informative signal about the true state $s^i \in \{b, r\}$. Each member can be of low or high ability. A high ability player receives a perfectly informative signal while a low ability player receives a correct signal with probability $\sigma \in (0.5, 1)$.⁷ Thus, there are four possible types of committee member $\{hb, hr, lb, lr\}$. With a slight abuse of terminology we will refer to $\{hb, hr\}$ as high types and $\{lb, lr\}$ as low types; $t^i \in \{h, l\}$, where $Pr(t^i = h) = q \in (0, 1)$. Abilities and signals are private information.

The principal gets a positive utility if the committee decision matches the state and zero otherwise. In [subsection 2.2](#), we examine the case where she cares more about one state than the other. The payoff of a committee member is simply the principal's posterior belief that he is of high ability, given by $\hat{q} \in [0, 1]$ - he gains no utility from the group decision per se. This is standard in models of career concerns (see Prat (2005); Fox and Van Weelden (2012); Levy (2007)).⁸

The timing of the game is as follows. Each committee member learns his ability and signal. Members can then communicate in a series of steps. First, in a simple straw poll, each member simultaneously announces a message $m^i \in \{m_b, m_r, m_\emptyset\}$, i.e. raises his hand in favor of B or R or abstains. Next, each member simultaneously announces a message $\tau^i \in \{\tau_h, \tau_l, \tau_\emptyset\}$, i.e. says he is a high type, low type or remains silent.⁹ After these two stages of communication, the committee has access to a coordination device - a publicly observable random draw from a uniform distribution $u[0, 1]$ which allows them to coordinate on a group decision. Finally, each member casts a vote $v^i \in \{v_B, v_R\}$, the group decision is taken and the true state is revealed to everyone.¹⁰

⁷We could allow for imperfect high type signals; the strategy of committee members in the most informative equilibria would remain largely unchanged. The only change is that there could now be a committee of two high types with conflicting signals. Under secrecy or mild transparency such a group would implement each decision with probability 0.5.

⁸This literature ignores the possible benefits to the principal of learning a member's ability level. One setting where the principal would not benefit is when labor markets are competitive and long term contracting is not possible (see discussion in Fox and Van Weelden (2012)).

⁹We could collapse all communication into one simultaneous round or could have the second stage of communication be sequential rather than simultaneous - the theoretical predictions would not change.

¹⁰We could also allow for abstention in the final voting stage but it would not change equilibrium predictions.

Once the group decision and state are revealed, the principal updates her beliefs on member abilities using all available information. Of course, the available information depends on the level of transparency. We compare three different regimes: secrecy, transparency and mild transparency. Under *secrecy* the principal only observes the group decision D . Under *transparency* she witnesses each member's messages m^i, τ^i , and final vote v^i . Under *mild transparency* she observes only the group decision and how each individual votes.

A committee member's strategy consists of a communication strategy and a voting strategy. A communication strategy is a pair $(\mathbf{m}^i, \boldsymbol{\tau}^i)$, where \mathbf{m}^i is a mapping from (s^i, t^i) into a probability distribution over messages $\{m_b, m_r, m_\emptyset\}$ and $\boldsymbol{\tau}^i$ is a mapping from (s^i, t^i) and messages exchanged in the straw poll into a probability distribution over announcements $\{\tau_h, \tau_l, \tau_\emptyset\}$. A voting strategy v^i is a mapping from signal, ability and messages exchanged in both communication rounds into a probability distribution over votes $\{v_B, v_R\}$.

We study symmetric perfect Bayesian equilibria under the three transparency regimes. As is standard in voting games, we restrict attention to strategies that are not weakly dominated.¹¹ As talk is cheap and payoffs depend on the principal's beliefs, there will be many equilibria in each of the three settings. We apply some restrictions to reduce the set of equilibria. Firstly, we restrict attention to equilibria in which h types vote in line with their private signal. Though other equilibria exist, it seems reasonable that h types will vote for the true state (especially as the principal knows they are perfectly informed). Our lab results support this restriction - h types vote to signal 98.2% of the time, rising to 100% in the final 5 periods.¹² Secondly, we follow the cheap talk literature (Crawford and Sobel (1982); Ottaviani and Sørensen (2001); Chen et al. (2008)) in focusing on the most informative equilibrium i.e. in each setting we look for equilibria where the maximal amount of information is shared across the two communication stages.¹³ Finally, we ignore equilibria with inverted language; for example where a message m_r is interpreted

¹¹In our setting this prevents cases where each player votes v_B simply because he expects the other player to do so, meaning no vote is pivotal.

¹²Alternatively, we could restrict attention to responsive and monotone strategies as in Fox and Van Weelden (2012).

¹³In the appendix we discuss the existence of other equilibria preferred by h types.

as “I have a signal in favor of state B ” or an announcement τ_h is interpreted as “I am of low ability”.

2.1. Equilibrium

Secrecy As the principal can only see the group decision, she must hold the same posterior \hat{q} for each individual. Committee members therefore have a common interest.

Proposition 1. *In the most informative equilibrium under secrecy, all members truthfully reveal their signal and ability, then jointly implement the decision with the highest posterior probability of matching the state. In case of a balanced posterior after two conflicting signals from low ability types, the committee coordinates on implementing each decision with probability 0.5. The probability of a group mistake in each state is $(1 - q)^2(1 - \sigma)$. Each member earns a payoff of $\frac{q}{1 - (1 - q)^2(1 - \sigma)}$ if the group decision matches the state, and zero otherwise.*

Proof. See Appendix □

It is hardly surprising that players with a common interest share their information. Coughlan (2000) shows that allowing communication between players with a common interest can lead to full information aggregation. There are other, less informative, equilibria in which players babble in one or both stages of communication. In the laboratory, however, Guarnaschelli et al. (2000) and Goeree and Yariv (2011) show that players with a common interest are overwhelmingly truthful.

The proposition states that a committee picks the decision most likely to match the state, and mixes with probability 0.5 in case of a balanced posterior belief. To see why this is the case, let all (lb, lr) committees implement B with (possibly degenerate) probability π . The expected utility of each committee member in equilibrium is then

$$\frac{0.5\pi q}{1 - (1 - q)^2(1 - \sigma)^2 - 2(1 - \pi)(1 - q)^2\sigma(1 - \sigma)} + \frac{0.5(1 - \pi)q}{1 - (1 - q)^2(1 - \sigma)^2 - 2\pi(1 - q)^2\sigma(1 - \sigma)}$$

where the first term corresponds to the expected evaluation in state B and the second to state R . The committees have to be indifferent between the two decisions in order to

mix, i.e. both terms have to be equal. This is only the case if $\pi = 0.5$ - when all (lb, lr) committees mix equally between implementing B and R .¹⁴ If (lb, lr) committees made one decision more often than another, the principal would take this into account - lowering his posterior belief following a correct decision in this state vis a vis the other state. Thus, the fact that committee mistakes are equally likely in both states stems entirely from the reputational concerns of l types.

Transparency Under this regime, the principal sees all stages of communication and observes each individual's vote.

Proposition 2. *In the most informative equilibrium under transparency, all members truthfully reveal their signal, information on abilities cannot be credibly communicated, and members vote according to their signal. The probability of a group mistake in state B is $(1 - q)^2(1 - \sigma)^2$, while the (much larger) probability of a mistake in state R is $1 - (q + (1 - q)\sigma)^2$. Each member earns a payoff of $\frac{q}{q+(1-q)\sigma}$ if their individual vote matches the state, and zero otherwise.*

Proof. See Appendix □

Here we cannot have full information sharing. If players were truthful about their ability with the principal watching, then l type players would receive a utility of zero. Instead, in an effort to appear competent, l types will mimic the strategy of h types so that in equilibrium no information on ability is revealed. This means information aggregation is incomplete when compared to secrecy: while signals are shared truthfully, players cannot differentiate the quality of those signals.

Even though some valuable information is shared, players actually ignore this when deciding how to vote. An lb player who sees a m_r message in the straw poll believes R is the most likely state. However, this lb player will be better off voting v_B in the final vote. The same holds for an lr player who sees a m_b message in the straw poll. Why is this? A h type would never vote against his signal, therefore any l type who switches his choice

¹⁴There are many ways they could achieve this, for example both voting v_B if the public coordination device gives a value below 0.5 and voting v_R otherwise.

between the straw poll and the final vote can be identified by the principal.¹⁵ So, l types will stick to their straw poll announcement in the final vote even when they believe it is wrong.

The combination of less information sharing and players “sticking to their guns” gives a higher aggregate probability of committee mistakes under transparency than under secrecy. However, there are differences across states. Transparency leads to far more mistakes in state R but somewhat fewer mistakes in state B . The reason is the different behavior of (lb, lr) committees under each regime. Under secrecy the committee will mix equally between implementing each decision, under transparency everyone votes according to their signal - meaning the status quo will always be upheld. If the true state is B , an (lb, lr) committee will never make a mistake under transparency while it will half of the time under secrecy.

Mild Transparency In this section, we look at an intermediate case of transparency. Here, the principal cannot observe any communication, but she does observe the individual votes of committee members as well as the final group decision. This corresponds to many real world cases where voting records are released but discussions are kept secret - as in the FOMC before 1993. Mild transparency also reflects the setting where the introduction of transparency leads to the emergence of pre-meetings. There members can discuss what to do before they are under the watchful eye of the principal. Alan Greenspan noted that introducing transparency in the FOMC would mean “a tendency would arise for one-on-one pre-meeting discussions, with public meetings merely announcing already agreed-upon positions or for each participant to enter the meeting with a final position not subject to the views of others” (quoted in Meade and Stasavage (2008)). Indeed, Swank and Visser (2013) have shown that pre-meetings can undo the benefits of transparency in their setting. In our setting, the following proposition shows that the most informative equilibrium under mild transparency gives the same outcomes as under secrecy.

Proposition 3. *In the most informative equilibrium under mild transparency communi-*

¹⁵This equilibrium is sustained by beliefs that any member who switches in the final vote is an l type. As nobody actually switches in equilibrium, this information set is never reached and so we are free to choose any beliefs of the principal off the equilibrium path.

cation, payoffs and the probability of mistakes are the same as under secrecy. In the voting stage, each member votes for the decision with the highest posterior probability of matching the state.

Proof. See Appendix □

The proposition implies that, if players do indeed play the most informative equilibrium, the emergence of pre-meetings would actually lead to more information aggregation and less mistakes than would otherwise be the case under transparency.¹⁶ There is a small difference between the secrecy and mild transparency equilibria - though it makes no difference to the probability of mistakes. Under secrecy only the group decision matters for belief updating - decisions in favor of B need not be unanimous. Under mild transparency, however, any member whose vote doesn't match the state will be revealed as an l type. Moreover, the principal knows that committee members perfectly share their information - so in a (h, l) committee both will vote correctly. The knock-on effect is that (lb, lr) committees must also vote unanimously. Otherwise both will be revealed as l types, even if one player votes correctly.

2.2. Optimal Level of Transparency

A natural question is which transparency regime the principal would prefer *ex ante*, assuming the most informative equilibria are played. As theory predicts the same outcome under secrecy and mild transparency, we compare the principal's utility under secrecy and transparency. In the analysis thus far, the principal's utility was irrelevant - committee members act to maximise the principal's beliefs of their ability, not to maximise her utility. To analyse the principal's welfare we must be more specific about her utility function than before. Let the principal gain a utility of x for a correct decision in state B and a utility of $1 - x$ in state R , where $x \in [0, 1]$, i.e. she may care more about correct decisions in one state than the other.¹⁷ The following proposition shows that the optimal level of transparency depends on x .

¹⁶There are other equilibria with less information sharing under mild transparency which are preferred by h types. See discussion in the appendix.

¹⁷This is equivalent to caring more about losses in one state than the other.

Proposition 4. *There always exists an $x^* > 0.5$ such that if $x < x^*$ the principal prefers secrecy while if $x > x^*$ she prefers transparency.*

Proof. See Appendix □

Why might transparency be preferred despite leading to more committee mistakes? By voting according to their signals, a transparent committee will seldom make the unanimous decision needed to implement R . As a result, they will make more mistakes than a secret committee in state R but less in state B . A principal with a high value of x is more wary of losses in state B , so will prefer transparency to secrecy.

Here, the fact that committee members ignore information under transparency actually helps a principal with a high x . Because the principal knows that members vote according to their signal, there is no pressure on committees to speak with one voice (a non-unanimous decision could come from a (h, l) or (l, l) committee). By ignoring information and "sticking to their signals", an (lb, lr) committee will always implement B . The same (lb, lr) committee under secrecy would implement B only 50% of the time. This is what gives lower mistakes in state B under transparency.

There are two points to note about the proposition: First, it is not something we can test in the lab - committee behavior does not depend on the value of x . Second, the threshold x^* is valid when there are no benefits to the principal of learning a committee member's type (for example due to perfectly competitive labor markets). If she did benefit from learning about types, then transparency - where it is easier to tell type apart - would be even more appealing to the principal.

3. Experiment

To test the main predictions of the model, we ran a laboratory experiment with three treatments - one for each level of transparency. In a slight departure from the model, we allowed for free-form communication in the second deliberation stage rather than simultaneous announcements and a public coordination device. The equilibria characterized in the previous section are unaffected by this change - the only sensible way subjects

could make use of the open chat is to share information on their ability and to coordinate. We wanted to test the theoretical predictions in this setting as (a) open deliberation is a feature of most real world committees, and (b) Goeree and Yariv’s (2011) results suggest that open chat may have strong effects on behavior. We decided against open chat in the first stage of deliberation as this might lead to very strange behavior under transparency.¹⁸

3.1. Experimental Design

We ran six sessions, two for each transparency regime (see Table 1). Each session consisted of 20 rounds. In the first round, subjects were randomly assigned to matching groups of nine people which remained fixed for all 20 rounds.¹⁹ Within each matching group, groups of three were randomly formed in every round. This was done to avoid the emergence of reciprocal behavior and at the same time provide a larger number of independent matching groups.²⁰ In each group, one person was randomly assigned the role of “observer” (the principal) and the other two the role of “voters” (committee members). Screenshots and instructions are in the online appendix. We set $\sigma = 0.55$ and $q = 0.25$ in all treatments.

Table 1: Experimental sessions

	N sessions	N matching groups	N subjects
Secrecy	2	5	45
Transparency	2	4	42
Mild Transparency	2	5	45

Note: All sessions were run at the DeSciL Lab at the ETH Zurich in May 2013 with 132 students (48% female) studying various majors at the ETH or the University of Zurich. Psychology students were not recruited.

¹⁸There is an incentive for l types to learn from the announcement of their partner. This might mean that, under transparency, a player who is slower in revealing his signal is interpreted as a l type by the principal. This, in turn, would lead to a “rush to announce” - where both players announce their signals immediately. This rushed open chat would then produce outcomes equivalent to the simultaneous announcements we use in the lab.

¹⁹In one session we had only 15 subjects and therefore only one matching group.

²⁰The matching groups are later treated as clusters in the estimation of standard errors.

There was a blue jar ($S = B$) and a red jar ($S = R$). The blue jar contained 11 blue and 9 red balls, the red jar contained 11 red and 9 blue balls. In each round, one jar was selected by a fair coin toss.²¹ A “well-informed voter” (h type) received a ball matching the color of the jar, an “informed voter” (l type) received a ball that was drawn from the jar.

Committee Members: On the first screen a member learns his type and the color of his ball, and then sends a message {red, blue, not specified} to the other member. On the next screen, he sees the message from his partner and has the opportunity to chat with him for 90 seconds.²² On the third screen, he can review the communication and make his final decision by voting for red or blue. On the final screen of each round, he learns the other member’s type, how the other member voted, the group decision, the true color of the jar, and his own payoff.

The Principal: On the first screen the player learns that she is a principal. Under secrecy she sees nothing else until voting is concluded. Once the committee has voted, she sees the true color of the jar and which group decision was taken. On this screen, she has to indicate her belief about the ability of a randomly chosen committee member, by entering a probability of him being a h type. Under mild transparency she also observes individual votes and has to evaluate each member separately. Under transparency she can see committee members’ messages and follow their chat in real time. After the vote she sees the same information as under mild transparency but can also review the communication (messages and chat) before evaluating each member. In all 3 treatments, on a final (feedback) screen she learns each member’s true type, their votes, the resulting group decision, and her own payoff.

Payoffs: The principal’s payoff for a correct group decision was 3 points if the jar was blue and 1 point if the jar was red. This corresponds to $x = 0.75$. In addition, the principal earned a number of points between 0 and 100 for accurate evaluation of the committee members’ types. Her points were determined by the following quadratic scoring rule.

²¹This has become the standard task in information aggregation experiments (e.g. Guarnaschelli et al. (2000); Battaglini et al. (2009); Goeree and Yariv (2011); Bhattacharya et al. (2014); Bouton et al. (2014)).

²²The timeout was not strictly enforced. When the time was up a message appeared on the screen asking them to finish their sentence and proceed. Most subjects did so immediately. The other few were asked to proceed after 120 seconds by an experimenter.

$$\text{Points} = \begin{cases} 100 - \frac{1}{100}(100 - Pr_j(t_i = h))^2 & \text{if committee member } i \text{ is an } h \text{ type and} \\ 100 - \frac{1}{100}(Pr_j(t_i = h))^2 & \text{if committee member } i \text{ is an } l \text{ type} \end{cases}$$

where $Pr_j(t_i = h)$ denotes principal j 's submitted probability (in per cent) that committee member i is a h type. This rule makes it optimal for expected payoff maximizing subjects to truthfully enter their beliefs (see, e.g., Nyarko and Schotter 2003) - participants were told as much in the instructions.²³ Under secrecy, the principal earned points from her single evaluation while under the other treatments one of her two evaluations was randomly chosen. We kept the principal's payoff from correct group decisions low relative to the payoff from accurate evaluations so as to limit the potential effects of social preferences.²⁴ Each committee member's payoff was determined by the principal. If she judged that member i had a $y\%$ chance of being a h type, that committee member would gain $2y$ points.

Four rounds were randomly chosen at the end of the session and the points earned in these rounds were converted to Swiss Francs at a rate of 1 point= CHF 0.15 (at the time of the experiment CHF 1 was worth USD 1.04). Subjects spent about 2 hours in the lab and earned CHF 47 on average, in addition to their show-up fee of CHF 10. Earnings per hour are comparable to an hourly wage for student jobs in Zurich.

3.2. Experimental Results

Aggregate Behavior

In the model we spoke of probabilities of a committee mistake; the analogue in the data are observed error rates, i.e. the share of committees that implement the wrong decision. In Table 2 we compare the observed error rates with equilibrium predictions. We see

²³More complicated belief elicitation procedures have been proposed for risk averse subjects (e.g., Offerman et al. (2009)). However, to avoid making the instructions overly complicated (and thus distracting subjects from the game) we chose to implement a standard quadratic scoring rule. We follow Schotter and Trevino (2014) in telling subjects that truthfully reporting their belief maximizes their expected payoff.

²⁴The concern here is that committee members would strive to make more correct decision in state B as these are more valuable to the principal.

immediately that the most informative equilibrium almost perfectly predicts error rates under secrecy. In the model, transparency generates more errors than secrecy in state R but less in state B . Indeed, this is borne out in the data.²⁵ Principals earned more points from group decisions under transparency than under secrecy - as the model with $x = 0.75$ predicts - though this difference was not statistically significant.²⁶

Table 2: Observed and Equilibrium Error Rates

True Color	$S = B$		$S = R$		Total	
	Observed	Equil.	Observed	Equil.	Observed	Equil.
<i>Secrecy</i>	28.3 (5.5)	25.3 [27.2]	25.8 (1.8)	25.3 [24.8]	27.0 (3.1)	25.3 [26.0]
<i>Transparency</i>	15.8 (1.0)	11.4 [8.2]	50.7 (6.4)	56.1 [61.9]	33.3 (3.5)	33.8 [33.9]
<i>Mild Transp.</i>	15.5 (1.2)	25.3 [24.0]	46.1 (3.6)	25.3 [25.3]	30.8 (1.8)	25.3 [24.7]

Note: Equil. = ex ante most informative equilibrium. Ex post equilibrium error rates, i.e. theoretical predictions after the realization of types and signals, are reported in brackets. Standard errors of the observed error rates (in parentheses) are adjusted for clustering in matching groups.

The most informative equilibrium predicts the same play under mild transparency and secrecy. Strikingly, however, error rates under mild transparency are statistically indistinguishable from those of transparency.²⁷ What explains these different error rates? To investigate, in Table 3, we analyse groups which “should” have differences across treatments, i.e. those with conflicting signals.

In (h, l) groups we expect no errors in state B under any treatment - the h type can unilaterally implement the correct decision. This is what we see in the data (apart from a single observation under mild transparency). In state R we expect no errors from conflicted (h, l) groups under secrecy, but a large number under transparency. We do find a significant difference between the two treatments, though the 44.8% error rate in transparency falls well short of the predicted 100%.²⁸ This suggests that, under transparency,

²⁵When compared to secrecy, the transparent committees performed significantly and substantially worse in state R at conventional levels (Wald-test, $p < 0.001$) but better in state B (Wald-test, $p = 0.024$).

²⁶ t -test, $p = 0.235$.

²⁷Wald-tests, $p = 0.526$ for $S = R$, and $p = 0.894$ for $S = B$.

²⁸In both Wald-tests we have $p < 0.001$

a significant share of l types are not voting according to their signal. In mild transparency the error rate in state R is significantly higher than the predicted value of zero, yet significantly lower than under transparency.²⁹ Clearly information is not being fully aggregated here.

Table 3: Error Rates in Groups with Conflicting Signals

		$S = B$		$S = R$	
		Observed	Equil.	Observed	Equil.
$\{h, l\}$ Groups	<i>Secrecy</i>	0	0	0	0
	<i>Transparency</i>	0	0	44.8 (8.3)	100
	<i>Mild Transp.</i>	3.3 (3.4)	0	22.6 (5.5)	0
$\{lb, lr\}$ Groups	<i>Secrecy</i>	38.7 (15.3)	50	54.3 (11.0)	50
	<i>Transparency</i>	27.5 (8.9)	0	83.3 (10.3)	100
	<i>Mild Transp.</i>	27.9 (6.8)	50	75.7 (4.4)	50

Note: Standard errors of the observed error rates (in parentheses) are adjusted for clustering in matching groups.

Next, we turn to (lb, lr) groups. Under secrecy we expect each decision to be implemented half of the time - otherwise committee evaluations would be lower in one state than the other. Indeed, the data is relatively close to an even split. In the model an (lb, lr) group under transparency would always implement B . Table 3 does show much higher error rates in state R than B but the difference between states is not as extreme as in theory. Some co-ordination on implementing R must be taking place here, though the accuracy is significantly lower than in (h, l) groups. Finally, observed error rates under mild transparency are virtually the same as those under transparency.³⁰

²⁹The first test has $p < 0.001$, the second $p = 0.026$.

³⁰Note that none of the differences in error rates between secrecy and the transparency treatments in these groups is statistically significant at the 5% level. In state R , however, the differences between secrecy and transparency ($p = 0.084$) and between secrecy and mild transparency ($p = 0.06$) are significant at the 10% level (Wald-tests). The differences in error rates between the two transparency regimes are not statistically significant even at the 10% level. The overall error rate under transparency is significantly higher in (lb, lr) than in (h, l) groups ($p < 0.001$).

Individual Behavior

Deliberation We begin our analysis of individual behavior with the first deliberation stage: the straw poll. Recall that in the most informative equilibrium all players truthfully reveal their signals - regardless of the transparency regime.

Table 4: (Non-)Truthful Straw Poll Messages from Different Types

	Type	Lying	Silent
<i>Secrecy</i>	<i>h</i> type	0	0
	<i>l</i> type	0.7 (.3)	1.8 (1.2)
<i>Transparency</i>	<i>h</i> type	1.4 (0.9)	0.7 (0.7)
	<i>l</i> type	8.3 (2.1)	10.5 (2.8)
<i>Mild Transp.</i>	<i>h</i> type	19.2 (6.6)	4.8 (2.0)
	<i>l</i> type	3.5 (1.1)	3.3 (1.4)

Note: Percentage of non-truthful messages (lying) and “not specified” messages (silent). Standard errors of the observed error rates (in parentheses) are adjusted for clustering in matching groups.

Table 4 shows two cases which clearly violate this: *l* types under transparency, and *h* types under mild transparency. In the first case, 8.3% of low types lie about their signal and another 10.5% stay silent.³¹ If an *l* type lies and then sticks with this lie, this would not be very costly - after all, σ -signals are not very informative. Why do some *l* types remain silent? Perhaps they hope to learn more about the true state and then vote accordingly. When we turn to the principals’ reactions, we will see whether this strategy pays off. In the second case, 19.2% of high types lie about their signal and another 4.8% stay silent.³² This behavior of *h* types is clearly inconsistent with the most informative equilibrium. It seems that some *h* types are hoarding their information in an attempt to separate from *l* types. Indeed, in the appendix, we show that there

³¹Both percentages are significantly larger than zero (Wald-tests, $p < 0.001$).

³²Both percentages are significantly larger than zero (Wald-tests, $p < 0.001$).

are equilibria of this type which yield higher payoffs for h types than full information sharing. Overall, however, the significant degree of lying is not enough to make the straw poll uninformative for l types.

Next, we turn to communication in open chat. After exchanging information about signals in the straw poll, the only relevant information left to discuss is member types.³³ We coded whether a type was announced in the chat, and if so, what type was announced. Under secrecy and mild transparency we expect truthful revelation of types. This is by and large what we observe in Table 5, though there is a small amount of lying under mild transparency. This suggests that the deviation of mild transparency from theory is explained by the frequent straw poll lies of h types rather than anything in the open chat.

Table 5: Chat Messages about Type

	Report Type	Claim		Lying
<i>Secrecy</i>	91.5 (1.5)	h type	25.9 (1.1)	0.7 (0.6)
		l type	74.1 (1.1)	0.5 (0.3)
<i>Transparency</i>	49.5 (8.7)	h type	80.1 (3.9)	51.4 (5.8)
		l type	19.9 (3.9)	1.8 (1.1)
<i>Mild Transp.</i>	89.0 (2.1)	h type	21.5 (1.0)	7.8 (5.3)
		l type	78.5 (1.0)	3.8 (1.4)

Note: The second column reports the fraction of committee members who report a type, the third the fraction of low and high type claims out of those reports and the fourth the fraction of lies out of these claims. Standard errors (in parentheses) are adjusted for clustering in matching groups.

We have no theoretical prediction for chat behavior in the transparency treatment - only that chat should be uninformative in equilibrium. In Table 5 we see that almost half of subjects announce a type - with 80.1% of those claiming to be a h type. This does not make communication informative per se, but a look at the truthfulness of these announcements shows that this stage is in fact informative. Of the 19.9% who announce an l type almost none are lying. In contrast, 51.4% of those claiming to be a h type are

³³We give examples of chats under secrecy and transparency in the online appendix.

lying. This rate of lying is actually lower than what we would expect if chat was pure babbling: given that h types only make up 25% of the population, 75% of h claims should be lies for the chat to be uninformative. What about those who don't announce a type? 89% of them are l types who possibly stay silent due to an aversion to lying or because they believe the principal would ignore the chat. In sum, a member's announcement (or lack thereof) is actually informative about his type.

Voting Do h types really vote as we assumed in our refinement? Yes, they vote according to their signal 98.2% of the time across all treatments.³⁴ Next, in Table 6, we look at how many l types vote against their signal when the other member announces a conflicting signal.

Under secrecy and mild transparency each member should announce their type to aid information aggregation. Indeed we saw overwhelmingly truthful revelation of types in Table 6, meaning that l types should believe those who claim to be h types. If one member claims to be a $hr(hb)$ type, an l member should always vote $v_R(v_B)$. We see that this is always the case under secrecy, and almost always under mild transparency. This is despite the moderate amount of deception under mild transparency. Under these two treatments we would expect (lb, lr) committees to implement each decision half of the time. Under mild transparency this involves players unanimously voting for each decision 50% of the time. Under secrecy players have a wider set of strategies, though all involve lb members switching at least 50% of the time and lr members switching at most 50% of the time.³⁵ In Table 6 we see that subjects get very close to 50% under secrecy and a little lower under mild transparency.

The really interesting case is transparency. The theory predicts that no player should ever switch their final vote from their straw poll announcement. In the lab, however, we see 30.8% of l types switching.³⁶ This is puzzling; switching from your straw poll

³⁴This is not statistically different from 100% (Wald-test, $p = 0.097$).

³⁵One strategy is the same as that of mild transparency - unanimously voting for each state 50% of the time. Another simple way to implement each decision 50% of the time is for the lr to vote to signal, and for the lb player to switch half of the time.

³⁶Table 6 reports voting against the signal and all subjects are predicted to truthfully report their signal in the straw poll. However, as we have seen some subjects do not report their signal truthfully, leading to slightly different switching rates between straw poll and final vote than between signal and final vote.

Table 6: Percentage of l Types Voting Against their Signal when Partner Reports a Conflicting Signal

Own signal blue		Other's Claim: h	Other's Claim: l	No Claim
<i>Secrecy</i>	Observed	100	42.4 (11.3)	50.0 (36.7) = 2 obs.
	Equil.	100	≥ 50	
<i>Transparency</i>	Observed	47.2 (5.9)	44.4 (9.4) = 9 obs.	37.5 (10.6)
	Equil.	0	0	
<i>Mild Transp.</i>	Observed	91.3 (4.5)	32.9 (1.9)	50.0 (12.2) = 6 obs.
	Equil.	100	50	
Own signal red		Other's Claim: h	Other's Claim: l	No Claim
<i>Secrecy</i>	Observed	100	45.9 (12.6)	40.0 (17.6) = 2 obs.
	Equil.	100	≤ 50	
<i>Transparency</i>	Observed	47.4 (5.1)	28.6 (7.9) = 2 obs.	57.1 (13.1) = 4 obs.
	Equil.	0	0	
<i>Mild Transp.</i>	Observed	95.23 (4.2)	32.4 (6.1)	34.8 (5.1)
	Equil.	100	50	

Note: Standard errors of the observed error rates (in parentheses) are adjusted for clustering in matching groups. Number of observations are reported if less than 10.

announcement should alert the principal that you are a l type, resulting in an evaluation of zero. In the next section we will investigate whether such switching was indeed punished.

Evaluations Next, we study how the principals evaluate. Table 7 shows that average evaluations are too high in all treatments and even significantly positive after wrong decisions.³⁷ Nonetheless, evaluations are much higher for correct decisions than incorrect ones, meaning the incentives to make a correct decision are about as strong as in theory.³⁸ In each treatment, h types receive higher evaluations on average than l types. Under

³⁷ t -tests, $p < 0.01$.

³⁸Evaluations look very similar if split by the state of the world (not reported here). There is no evidence that principals reward correct decisions in state B more than in state R .

transparency this is even the case when both types vote correctly, showing that principals can distinguish the two types well.

Table 7: Evaluations

		Evaluation		
		Avg.	Dec. Corr.	Dec. Wr.
<i>Secrecy</i> (group decision)	Equil.	25	33.5	0
	Observed	41.6 (1.5)	54.5 (3.2)	6.9 (1.9)
	Observed <i>h</i> types	53.5 (3.4)	53.5 (3.4)	
	Observed <i>l</i> types	37.4 (1.5)	55.0 (3.3)	6.9 (1.9)
<i>Transparency</i> (individual decision)	Equil.	25	37.7	0
	Observed	37.2 (3.0)	48.2 (1.7)	14.7 (4.0)
	Observed <i>h</i> types	60.8 (3.2)	62.8 (3.8)	
	Observed <i>l</i> types	29.4 (2.9)	40.2 (2.2)	14.6 (4.0)
<i>Mild Transp.</i> (individual decision)	Equil.	25	33.5	0
	Observed	35.7 (1.9)	47.1 (1.9)	11.7 (1.6)
	Observed <i>h</i> types	47.7 (1.2)	48.3 (1.1)	
	Observed <i>l</i> types	31.9 (2.3)	46.4 (2.6)	11.8 (1.6)

Note: Standard errors (in parentheses) are adjusted for clustering in matching groups.

In Table 8 (models 1, 3 and 5) we regress the principal's evaluations on the information she sees before evaluating. Making the wrong group decision greatly reduces a member's evaluation under secrecy, but not under the other two treatments (where individual voting data is observed). If the most informative equilibrium was played under mild transparency, then the group decision should matter (as the principal expects members to share information). The fact that only individual decisions matter here means that *h* types face no risk in lying and leading *l* types astray.

Under transparency, anything apart from a correct message and correct individual vote results in a lower evaluation - so switchers are actually punished. Interestingly, claiming

Table 8: Evaluation Responses

	1	2	3	4	5	6
	Sec	Sec-Lin	Trans	Trans-Lin	Mild	Mild-Lin
Group decision D wrong	-47.6*** (4.2)	-36.1*** (1.7)	0.7 (2.9)	-1.6 (2.9)	-0.9 (3.3)	-2.3 (4.4)
Individual vote v wrong					-34.8*** (2.1)	-32.7*** (3.0)
<i>Combinations of m and v, reference category: message and vote are correct</i>						
v wrong, m wrong			-34.0*** (4.9)	-34.4*** (3.5)		
v wrong, m right			-31.3*** (3.9)	-19.3 (10.1)		
v right, m wrong			-20.6** (3.8)	-27.9*** (3.0)		
v right, silent in straw poll			-15.5** (4.6)	-30.3*** (3.7)		
v wrong, silent in straw poll			-40.6*** (5.2)	-32.3*** (3.0)		
Claimed to be type h in chat			15.3*** (2.1)	29.7*** (2.4)		
Claimed to be type l in chat			-17.5* (5.9)	-7.5 (3.8)		
Constant	54.5*** (3.4)	36.1*** (1.7)	46.2*** (2.2)	28.2*** (3.6)	47.1*** (1.9)	35.5*** (2.5)
N observations	600	600	560	560	600	600
N clusters	5	5	4	4	5	5
R^2	0.49	0.13	0.29	0.33	0.31	0.14

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors (in parentheses) are adjusted for clusters in matching groups. The dependent variable in the comparison columns 2,4 and 6 is rescaled to the same range as the dependent variable in 1,3 and 5 and takes the value 100 if the subject is a h type and 0 if not.

to be a h type leads to a higher evaluation than remaining silent while claiming to be an l type gives a lower evaluation. It seems the principal was aware of the relatively high level of truth-telling seen in Table 5.

(Best) Responses In the previous sections we saw that the observed behavior of both principals and committee members deviated somewhat from theoretical predictions. But perhaps committee members take the principals deviation into account when deciding on their own strategy. Similarly, principals may anticipate that committee members will not behave in line with theory. Therefore, we now examine whether individuals are best responding to the behavior of other players.

Principals: The evaluations that principals give are too high to maximise their payoffs. This may be because they find it difficult to update beliefs correctly. They do perform better over time. Table OA1 in the online appendix shows that principals lower their evaluations in later rounds, which suggests learning. However, committee members who vote incorrectly often receive positive evaluations - even in the final rounds of play. These positive evaluations occur when updating is extremely simple and an evaluation of zero would maximise the principal's payoff. This suggests that principals have social preferences towards committee members as their evaluations directly affect members' payoffs. Another factor that might play a role is that the scoring rule does not correct for risk-aversion. Entering 50% as an evaluation gives you a certain payoff of 75 whereas any other evaluation leaves you with a lottery if you attach positive probabilities to either type. Risk-averse subjects would therefore be drawn toward 50% which is double the ex ante probability of a high type q .

However, we are primarily interested in how evaluations differ across regimes and player behavior. Does the principal react in the right direction to each piece of information? To answer this we can compare columns 1, 3 and 5 to columns 2, 4 and 6 in Table 8. The odd columns have the principals' actual evaluations as the dependent variable while the even columns are linear probability models where the dependent variable is binary, taking a value of 100(%) if the member is an h type and 0 if not. If principals were best responding they would behave as in the even columns. We see that, in fact, they

come very close. Principals use all pieces of information in the right direction, often with accurate magnitudes.

Committee Members: Under secrecy committee members best respond to the principal's behavior, as well as the behavior of other committee members. They first share their private information in the group and then implement the decision most likely to match the state. We know that behavior under mild transparency deviates from theory, but are committee members best responding given these deviations? High types vote to signal - as they should. Some tell the truth while others lie about their signal. Are some of these players doing better than others? Actually, as realised evaluations do not depend on group decisions, any level of information sharing by h types is a best response. Most h types do reveal their information. This makes it optimal for l types to switch their vote when they encounter a h type with a conflicting signal. This is exactly what we observe in Table 6 - l types are best responding under secrecy and mild transparency.

Under transparency, the vast majority of h types best respond by announcing their signal and type truthfully and then voting for the correct decision. However, a substantial fraction of low types do not best respond. In the most informative equilibrium, an l type should truthfully announce his signal and stick to this signal in the final vote. Whichever type he claims to be, this should be uninformative in equilibrium. The following analysis shows that it is in fact best for l types in the lab to (i) truthfully announce their signal, (ii) claim to be a h type, and (iii) vote to signal.

In the straw poll, 8.3% of l types lie, while 10.5% remain silent. Neither is a best response. Staying silent here may be part of a learning strategy whereby an l type waits to see the signal of his partner. If the principal believed the straw poll was idle babbling, then this might pay off. However, from Table 8, we see that this strategy results in lower evaluations than announcing the signal truthfully.

In the open chat many l types either truthfully reveal their type, or simply avoid declaring a type. This leads to lower evaluations than those who claim to be h types. The number of committee members who claim to be an l type goes down from 12.9% in the first ten to 6.8% in the last ten rounds, while the number remaining silent climbs from 45.7%

to 55.3%.³⁹ There are several explanations for the behavior of these l types. They may believe that the principal will ignore announcements on type as pure babbling. It could be that some l types are trying to achieve the correct group decision, so don't want to fool their partner by pretending to be a h type. Or, it may be that some players have an aversion to lying (Gneezy et al. (2013) and Fischbacher and Föllmi-Heusi (2013)).

In theory we would never expect an l type to switch between his straw poll announcement and his final vote. Doing so would perfectly reveal his type. In the lab we actually see a large amount of switching from l types. The fact that (i) evaluators give too much credit for correct votes after incorrect messages, and (ii) around half the claims of being a h type are legitimate, makes it not very costly for an l type to switch when his partner claims to be a h type with a conflicting signal.⁴⁰ Those who switch do not play a best response but it only costs them CHF 0.65 on average. Switching if the other member does not claim to be a h type is more costly and as a consequence occurs less often (see Table 6).⁴¹

4. Discussion

In this paper we constructed a model of committee decision-making in which members are career concerned. We examined how the incentives of committee members to share their private information and vote strategically varies with the level of transparency, i.e. how much of the decision process the principal observes. We show that transparency harms information aggregation and leads to more committee mistakes than secrecy.

There are a number of reasons to question the predictive power of the model. Firstly, the different outcomes are driven by subtle strategic behavior of “low type” committee members. These range from strategic communication, to strategic voting and coordination of votes among group members. However, a number of experimental studies have

³⁹Both of these differences are statistically significant with $p = 0.031$ for the decrease in l claims and $p < 0.001$ for the increase in no claims.

⁴⁰The difference in evaluations is 4.3 percentage points and is not statistically different from zero (t -test, $p = 0.399$).

⁴¹The difference in evaluations is 9.9 percentage points (1.48 CHF) and is not statistically different from zero (t -test, $p = 0.409$) as the number of observations is low for this scenario.

shown for other settings that some subjects communicate and vote far less strategically than theory predicts. For communication this appears to be especially true when players are allowed to chat freely (as they would in real life committees). Secondly, other experiments have shown that at least some people exhibit psychological costs of lying which casts some doubt on whether committee members will really behave only to maximise their reputations. Thirdly, the predictions of the model are based on an equilibrium refinement - the *most informative equilibrium*. This refinement is common in cheap talk games, though whether it is a good predictor of behavior is an open question (see Sobel (2013)). Given these concerns, we were eager to test our model empirically. By bringing it to the laboratory, we gained a number of insights.

The main predictions of the model are borne out in the lab: there is less information sharing and more committee mistakes under transparency than under secrecy. The difference in outcomes is in fact driven by players understanding the subtle incentives to maximise their expected payoffs. Under secrecy, committees with a balanced posterior appear to realise that they should coordinate to implement each decision with equal probability. Under transparency, 87.2% of low types understand that they should not reveal their true type as this will hurt their evaluation. Also under transparency, a majority of low types vote in line with their private signal even when the evidence points in favor of the other decision. They do this as voting against their signal announcement would out them as a low type.

The difference in mistakes between transparency and secrecy is large but not as pronounced as the theory predicts. This is largely because some committee members are more truthful under transparency than in the most informative equilibrium. While it might seem surprising that some players tell the truth when it is optimal to lie, this is in line with previous experimental results showing that many people face a psychological cost of lying. This aversion to lying suggests that the negative effect of transparency on group decision-making, though important, may not be as severe as in theory.

The predictive performance of the most informative equilibrium is mixed: it performs very well under secrecy, quite well under transparency, and not particularly well under mild transparency - where high types are much less truthful than predicted. As we show

in the appendix, the most informative equilibrium under secrecy and transparency coincides with the best equilibrium for high ability players (those with the most information to share). Under mild transparency, however, the most informative equilibrium is payoff-dominated for high ability types; they prefer equilibria in which no information is shared. It is perhaps unsurprising, therefore, that there is a substantial degree of deception from high ability types under mild transparency. As a result, group errors and player evaluations here are much closer to the transparency case than to secrecy, casting doubt on the predictive power of the most informative equilibrium. This is a finding which is relevant beyond this particular context. Sobel (2013) called for experimental tests of the predictive power of the most informative equilibrium in cheap talk games in which other equilibria exist which are preferred by at least one type of player.

Overall, despite some deviations from theoretical predictions, we find large differences in behavior between the three regimes: information aggregation is successful under secrecy while it breaks down, for different reasons, under transparency and mild transparency. This results in fewer mistakes under secrecy than the two other treatments and suggests that the level of transparency is a highly important element of institutional design - setting it wrong might indeed have considerable negative consequences for a principal. The difference in the level of truth-telling between treatments contrasts sharply with the results of Goeree and Yariv (2011). They find that free-form communication greatly diminishes institutional differences; players are overwhelmingly truthful regardless of the voting rule. Our results show that the degree of transparency in a committee strongly affects members' behavior even with open chat. This suggests that the level of openness may be a more important institutional choice than that of the voting rule.

References

- BATTAGLINI, M. AND U. MAKAROV (2014): “Cheap talk with multiple audiences: An experimental analysis,” *Games and Economic Behavior*, 83, 147–164.
- BATTAGLINI, M., R. B. MORTON, AND T. R. PALFREY (2009): “The Swing Voter’s Curse in the Laboratory,” *Review of Economic Studies*, 77, 61–89.
- BENTHAM, J. ([1816] 1999): *Political Tactics*, Oxford: Oxford University Press.
- BHATTACHARYA, S., J. DUFFY, AND S.-T. KIM (2014): “Compulsory versus voluntary voting: An experimental study,” *Games and Economic Behavior*, 84, 111–131.
- BOUTON, L., M. CASTANHEIRA, AND A. LLORENTE-SAGUER (2014): “Divided Majority and Information Aggregation: Theory and Experiment,” .
- CAI, H. AND J. T.-Y. WANG (2006): “Overcommunication in strategic information transmission games,” *Games and Economic Behavior*, 56, 7–36.
- CHEN, Y., N. KARTIK, AND J. SOBEL (2008): “Selecting Cheap-Talk Equilibria,” *Econometrica*, 76, 117–136.
- COHEN, J. (1996): “Deliberation and Democratic Legitimacy,” in *Deliberative Democracy: Essays on Reason and Politics*, ed. by J. Bohman and W. Rehg, Cambridge, MA: MIT Press, 67–92.
- COUGHLAN, P. (2000): “In Defense of Unanimous Jury Verdicts : Mistrials , Communication , and Strategic Voting,” *American Political Science Review*, 94, 375–393.
- CRAWFORD, V. AND J. SOBEL (1982): “Strategic information transmission,” *Econometrica*.
- CROSS, J. P. (2013): “Striking a pose: Transparency and position taking in the Council of the European Union,” *European Journal of Political Research*, 52, 291–315.
- FISCHBACHER, U. AND F. FÖLLMI-HEUSI (2013): “Lies in Disguise-an Experimental Study on Cheating,” *Journal of the European Economic Association*, 11, 525–547.

- FOX, J. AND R. VAN WEELDEN (2012): “Costly Transparency,” *Journal of Public Economics*, 96, 142–150.
- GERSBACH, H. AND V. HAHN (2008): “Should the Individual Voting Records of Central Bankers be Published?” *Social Choice and Welfare*, 30, 655–683.
- GNEEZY, U. (2005): “Deception: The role of consequences,” *The American Economic Review*, 95, 384–394.
- GNEEZY, U., B. ROCKENBACH, AND M. SERRA-GARCIA (2013): “Measuring lying aversion,” *Journal of Economic Behavior & Organization*, 93, 293–300.
- GOEREE, J. K. AND L. YARIV (2011): “An Experimental Study of Collective Deliberation,” *Econometrica*, 79, 893–921.
- GUARNASCHELLI, S., R. D. MCKELVEY, AND T. R. PALFREY (2000): “An Experimental Study of Jury Decision Rules,” *American Political Science Review*, 94, 407.
- HABERMAS, J. (1996): *Between Facts and Norms*, Cambridge, MA: MIT Press.
- LANDA, D. AND A. MEIROWITZ (2009): “Game Theory, Information, and Deliberative Democracy,” *American Journal of Political Science*, 53, 427–444.
- LEVY, G. (2007): “Decision Making in Committees: Transparency, Reputation, and Voting Rules,” *American Economic Review*, 97, 150–168.
- MALENKO, N. (2013): “Communication and decision-making in corporate boards,” *Review of Financial Studies*.
- MEADE, E. E. (2005): “The FOMC: Preferences, Voting, and Consensus,” 93–102.
- MEADE, E. E. AND D. STASAVAGE (2008): “Publicity of Debate and the Incentive to Dissent: Evidence from the US Federal Reserve,” *The Economic Journal*, 118, 695–717.
- NYARKO, Y. AND A. SCHOTTER (2003): “An experimental study of belief learning using elicited beliefs,” *Econometrica*, 70, 971–1005.

- OFFERMAN, T., J. SONNEMANS, G. VAN DE KUILEN, AND P. P. WAKKER (2009): “A Truth Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes,” *Review of Economic Studies*, 76, 1461–1489.
- OTTAVIANI, M. AND P. SØRENSEN (2001): “Information aggregation in debate: who should speak first?” *Journal of Public Economics*, 81, 393–421.
- PRAT, A. (2005): “The Wrong Kind of Transparency,” *American Economic Review*, 95, 862–877.
- SCHOTTER, A. AND I. TREVINO (2014): “Belief Elicitation in the Laboratory,” *Annual Review of Economics*, 6, 103–128.
- SOBEL, J. (2013): “Ten possible experiments on communication and deception,” *Journal of Economic Behavior & Organization*, 93, 408–413.
- STASAVAGE, D. (2007): “Polarization and publicity: rethinking the benefits of deliberative democracy,” *Journal of Politics*, 69, 59–72.
- SWANK, J., O. H. SWANK, AND B. VISSER (2008): “How Committees of Experts Interact with the Outside World: Some Theory, and Evidence from the FOMC,” *Journal of the European Economic Association*, 6, 478–486.
- SWANK, O. H. AND B. VISSER (2013): “Is Transparency to No Avail?” *Scandinavian Journal of Economics*, 115, 967–994.
- VISSER, B. AND O. H. SWANK (2007): “On Committees of Experts,” *Quarterly Journal of Economics*, 337–372.
- WANG, J. T.-Y., M. SPEZIO, AND C. F. CAMERER (2010): “Pinocchio’s Pupil: Using Eyetracking and Pupil Dilation to Understand Truth Telling and Deception in Sender-Receiver Games,” *American Economic Review*, 100, 984–1007.
- WILLIAMS, J. C. (2012): “Opening the Temple,” *Federal Reserve Bank of San Francisco 2011 Annual Report*.

Appendix

Proof of Proposition 1

The most informative outcome of the two communication stages is, by definition, where each player reveals his signal and ability. Let a candidate for equilibrium be where members truthfully announce their signal and ability, then implement the decision with the highest posterior probability of matching the state. In case of a balanced posterior (after two conflicting signals from low ability types), the committee coordinates on implementing each decision with probability 0.5. Assuming this is an equilibrium, the probability of mistakes and member evaluations are as follows.

Mistakes: A mistake occurs in state B when we have a pair (lr, lr) (with probability $(1 - q)^2(1 - \sigma)^2$) or a pair (lb, lr) who implement R (with probability $(1 - q)^2(1 - \sigma)\sigma$). Thus the total probability of implementing decision R in state B is $(1 - q)^2(1 - \sigma)$. The case of a mistake in state R is symmetric.

Evaluations: Given truthful communication, an incorrect group decision reveals that there are no h types on the committee; thus each player gets an evaluation of zero. Instead, when a committee makes the correct decision the principal updates her beliefs in the following way:

$$\begin{aligned} \hat{q}_{sec}(D = S) &= \frac{\sum_{k=0}^2 \frac{k}{2} (1 - Pr_{sec}(D \neq S | k \# \text{ of } h \text{ types})) \binom{2}{k} q^k (1 - q)^{2-k}}{1 - Pr_{sec}(D \neq S)} \quad (1) \\ &= \frac{q^2 + 0.5q(1 - q)2}{1 - (1 - q)^2(1 - \sigma)} \\ &= \frac{q}{1 - (1 - q)^2(1 - \sigma)} \end{aligned}$$

We will show this is indeed an equilibrium and discuss its uniqueness.

Existence: As h types always vote according to their signal, we have that $\hat{q}_{sec}(D = S) > \hat{q}_{sec}(D \neq S) = 0$. Thus each member strictly prefers the committee decision to match the state. As the principal only observes the group outcome, a player would only

have an incentive to deviate if it increases the probability of a correct group decision. Any group with a h member will implement the correct decision, so in these groups no player has an incentive to deviate. It remains to show that there is no incentive for any member to deviate in (l, l) groups. In an (lr, lr) group implementing R would give an expected evaluation of $\frac{\sigma^2}{\sigma^2+(1-\sigma)^2}\hat{q}_{sec}(D = S)$, while implementing B would lead to a lower expected evaluation of $\frac{(1-\sigma)^2}{\sigma^2+(1-\sigma)^2}\hat{q}_{sec}(D = S)$. Sending a non-truthful message can only lower the expected evaluation (for example by inducing the other lr member to vote v_B). Hence, there is no incentive for any member here to deviate. Similarly, there is no profitable deviation in a (lb, lb) group. In a (lb, lr) group the proposed equilibrium has members implementing each decision with probability 0.5. This yields an expected payoff for each member of $\frac{0.5q}{1-(1-q)^2(1-\sigma)^2-(1-q)^2\sigma(1-\sigma)}$. As the group posterior is balanced and all (lb, lr) groups are mixing equally between each decision, the expected evaluation from implementing B or R is the same. Therefore there is no strict incentive to deviate from this mixing strategy. As there is no benefit in implementing one decision more than another, there is also no incentive to deviate from truth-telling in communication. Thus, the candidate equilibrium is in fact an equilibrium.

Uniqueness: Restricting attention to full information sharing, there is still more than one equilibrium. For example, when the group's posterior favours B it could be that one or both members vote v_B . Furthermore, in a (lb, lr) group, there are many voting strategies which will implement each group decision with probability 0.5. Nonetheless, in any equilibrium with full information sharing the group will implement the decision with the highest posterior probability of matching the state, and will implement each decision with probability 0.5 if they have a balanced posterior. To see why mixing with probability 0.5 is unique, let all (lb, lr) committees implement B with (possibly degenerate) probability π . The expected utility of each committee member in equilibrium is then

$$\frac{0.5\pi q}{1-(1-q)^2(1-\sigma)^2-2(1-\pi)(1-q)^2\sigma(1-\sigma)} + \frac{0.5(1-\pi)q}{1-(1-q)^2(1-\sigma)^2-2\pi(1-q)^2\sigma(1-\sigma)}$$

where the first term corresponds to the expected evaluation in state B and the second to state R . To be willing to mix, the committees must be indifferent between the two

decisions, i.e. both terms have to be equal. This is only the case if $\pi = 0.5$.

Proof of Proposition 2

First we show that there can be no information on ability revealed in communication. Suppose there existed an equilibrium in which each player truthfully reveals his signal and ability. As the principal sees these announcements, l types would get an evaluation $\hat{q} = 0$ while h types would get an evaluation $\hat{q} = 1$. An lb type would have an incentive to deviate, announcing τ_h rather than τ_l and thus pooling with hb types. This deviation would earn him an evaluation $\hat{q} = 1$ if the state of the world is B , a clear improvement on zero. An lr type has the same incentive to pool with hr types by announcing τ_h . Therefore, full truth-telling cannot be an equilibrium under transparency. Can any information on ability be revealed in communication? No. Fix arbitrary communication and voting strategies for hr and hb types. As the principal observes all individual behavior under transparency, only actions consistent with the strategies of h types will receive positive evaluations. As such, lb and lr types will have an incentive to perfectly mimic the strategies of hb and hr types. By matching high types' distribution over ability announcements, low types succeed in signal jamming - no information on ability is revealed.

The most informative outcome of the two communication stages is, therefore, where each player reveals his signal truthfully and where no meaningful information on ability is communicated. Let a candidate for equilibrium be where members truthfully announce their signal, hb and lb members follow the same strategy in announcing ability, τ^b , hr and lr members follow the same strategy in announcing ability, τ^r , and each member votes according to his signal. Assuming this is an equilibrium, the probability of mistakes and member evaluations are as follows.

Mistakes: In state B a mistake only occurs when we have an (lr, lr) committee, as R is wrongly implemented. This occurs with probability $(1 - q)^2(1 - \sigma)^2$. In state R a correct decision will be made by a committee composed of hr or lr members. Thus, the probability of a mistake is $1 - (q + (1 - q)\sigma)^2$.

Evaluations: Only members who receive a signal in line with the state will be given

positive evaluations. All h types will receive the correct signal as will a share σ of l types. The principal will thus give an evaluation $\frac{q}{q+(1-q)\sigma}$ if a member announces the correct signal and also votes for that signal.

We will show this is indeed an equilibrium and discuss its uniqueness.

Existence: In this candidate for equilibrium, each member's message m^i and vote v^i always agree. Furthermore, hb and lb players only make ability announcements permitted in the strategy τ^b , while hr and lr players only make ability announcements permitted in the strategy τ^r . To check the existence of this equilibrium we need to fix the principal's off-path beliefs. We set off-path beliefs to be that the only positive evaluations are $\hat{q}(m_b, \tau^i, v_B | S = B, \tau^i \in \tau^b)$ and $\hat{q}(m_r, \tau^i, v_R | S = R, \tau^i \in \tau^r)$. Given these beliefs, it is immediate that hb and hr members have no incentive to deviate. An lb member knows that he must announce either m_b or m_r to gain a positive evaluation. As his prior is informative, his best response is to announce m_b . Given the beliefs of the principal, the only course of action open to this lb member is then to announce an ability consistent with τ^b , and then vote v_B . Any other course of action will lead to an evaluation of zero for sure.

Uniqueness: The strategies τ^b and τ^r are arbitrary. All that an equilibrium requires is that hb and lb follow the same communication strategy, as do hr and lr . For example, it could be that all members announce τ_h or all remain silent or perhaps $\tau^b = \tau_h$ while $\tau^r = \tau_l$. Nonetheless, in each equilibrium there is no information on ability levels revealed in communication. In the set of most informative equilibria, only equilibria in which all players truthfully announce their signal and then vote according to their signal exist. To see this, note that h types will vote according to their signal (by assumption), so the only possible deviation in this set is for l types not to vote according to their signal. As the messages and votes of h types always agree, an l type (who truthfully reports his signal) would reveal himself as an l type if he voted against his signal, getting an evaluation of zero.

Proof of Proposition 3

The most informative outcome of the two communication stages is, by definition, where each player reveals his signal and ability. Let a candidate for equilibrium be where mem-

bers truthfully announce their signal and ability, then unanimously vote for the decision with the highest posterior probability of matching the state. In case of a balanced posterior (after two conflicting signals from low ability types), the committee unanimously votes for each decision with probability 0.5. Assuming this is an equilibrium, the probability of mistakes and member evaluations are as in Proposition 1. We will show this is indeed an equilibrium and discuss its uniqueness.

Existence: In this candidate for equilibrium, all group decisions are unanimous. To check its existence we must set off-path beliefs for non-unanimous decisions. We impose that $\hat{q}_i(v_B^i, v_B^j | S = B) > \hat{q}_i(v_B^i, v_R^j | S = B)$, $\hat{q}_i(v_B^i, v_R^j | S = R)$. Given these beliefs, it is immediate that a h type will never have an incentive to deviate from the proposed strategy. Similarly, the off-path beliefs make it optimal for those in a (lr, lr) or (lb, lb) group to tell the truth, while the fact that $\hat{q}_i(v_B^i, v_B^j | S = B) = \hat{q}_i(v_R^i, v_R^j | S = R) > \hat{q}_i(v_B^i, v_B^j | S = R) = \hat{q}_i(v_R^i, v_R^j | S = B) = 0$ means that there is no incentive to deviate from a strategy which unanimously implements the decision with the highest posterior probability of matching the state. In a (lb, lr) group the proposed equilibrium has members unanimously implementing each decision with probability 0.5.⁴² As shown in the proof of Proposition 1, there is no incentive for any member to deviate from this equal mixing strategy. The requirement that (lb, lr) groups vote unanimously is sustained by the off-path beliefs. Thus, the candidate equilibrium is in fact an equilibrium.

(ii) Uniqueness: This is the only equilibrium under mild transparency in which all members fully share their information, and h types vote according to their signal. This uniqueness is guaranteed by the principal's off-path beliefs, which lead l types to vote unanimously for the decision most likely to match the state. If the principal's off-path beliefs were instead $\hat{q}_i(v_B^i, v_B^j | S = B) \leq \hat{q}_i(v_B^i, v_R^j | S = B)$ then truthful communication would break down. For more on this see below.

⁴²This is achieved using the coordination device (e.g both voting v_B if the draw is below 0.5 and voting v_R if it is above).

Proof of Proposition 4

The principal will be better off under transparency than under secrecy when

$$\begin{aligned} x[Pr_{tran}(D = B|B)] + (1-x)[Pr_{tran}(D = R|R)] > \\ x[Pr_{sec}(D = B|B)] + (1-x)[Pr_{sec}(D = R|R)] \end{aligned}$$

which can be rearranged to

$$(1-x)[Pr_{tran}(D = B|R) - Pr_{sec}(D = B|R)] < x[Pr_{sec}(D = R|B) - Pr_{tran}(D = R|B)]$$

substituting in with the values from proposition 1 and 2 and rearranging we get

$$\begin{aligned} \frac{(1-q)^2(1-\sigma)\sigma + 2q(1-q)(1-\sigma)}{2(1-q)^2(1-\sigma)\sigma + 2q(1-q)(1-\sigma)} < x \\ \frac{(1-q)\sigma + 2q}{2(1-q)\sigma + 2q} \equiv x^* < x \end{aligned} \quad (2)$$

Thus, secrecy is preferred if $x < x^*$ while transparency is preferred if $x > x^*$.

When is the most informative equilibrium payoff dominated?

In the most informative equilibrium, a h player will have a higher expected utility under transparency than secrecy or mild transparency:

$$\frac{q}{1 - (1-q)(1-\sigma)} > \frac{q}{1 - (1-q)^2(1-\sigma)} \quad (3)$$

An l type player must weigh these expected evaluations by the probability of voting for the correct state. Unsurprisingly, l types have a higher expected utility when they can pool with h types - they prefer secrecy to transparency:

$$\frac{\sigma q}{1 - (1-q)(1-\sigma)} < \frac{((1-q)\sigma + q)q}{1 - (1-q)^2(1-\sigma)} \quad (4)$$

As h types are those with the bulk of information to share, we examine if the most informative equilibrium in each setting is the equilibrium with the highest payoff for h types. Here, we only relax the assumption that the most informative equilibrium is played. We maintain our refinement assumptions regarding voting behavior.

Secrecy Here players face a common-value problem. The most informative equilibrium allows players to aggregate their private information and then make a decision which maximizes the group (and each player's) expected evaluation. No player can be better off by withholding information.

Transparency Here the most informative equilibrium involves all players voting to signal. The combination of voting to signal and individual votes being observed means the principal can distinguish h and l types very well. In fact, when all players vote to signal, h types get the highest evaluation they can get in any equilibrium.

Mild Transparency Here the most informative equilibrium coincides with that of secrecy. However, as the principal now observes individual votes, a h type can achieve a higher payoff in another equilibrium in which he separates from l types. That is, there are equilibria with less information sharing (or none) which payoff dominate the most informative equilibrium for h types. One such case is where no information is credibly communicated; for example, h types may mix between announcing m_b and m_r with equal probability. With no information communicated, the best response of l types is to vote to signal. In this polar opposite to the most informative equilibrium we see that h type players gain a higher payoff, $\frac{q}{1-(1-q)(1-\sigma)}$. There are a series of equilibria between these two poles which are preferred by h types to full truth-telling. In these equilibria, all players reveal their ability, however h types mix between truthfully revealing their signal and remaining silent while l types vote against their private signal when they see a conflicting message from a h type. Unlike the truth-telling case, in all these "preferred equilibria" the payoff of each committee member is independent of his partner's action. Indeed, as table 10 shows, this is what we find in our laboratory setting.

Online Appendix

OA1 Evaluations over time

Table OA1: Evaluations over rounds

MOA1		
<i>dependent variable: evaluation</i>		
Round number	-0.36**	(0.12)
<i>Secrecy</i>		
decision correct	58.21***	(3.75)
decision wrong	10.62***	(1.78)
<i>Transparency</i>		
decision correct	51.89***	(2.05)
decision wrong	18.58***	(4.14)
<i>Mild Transparency</i>		
decision correct	50.86***	(2.60)
decision wrong	15.36***	(2.39)
<i>N</i> observations	1760	
<i>N</i> clusters	14	
R^2	0.31	

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses are adjusted for clusters in matching groups. The model was estimated without a constant. The R^2 statistic was computed with a constant, leaving out one group as a reference group.

OA2 Sample chat conversations (translated from German)

Further chat protocols are available from the authors upon request.

Conversation 1 (Secrecy):

hey (voter 1)

type h (voter 2)

hello (voter 2)

I am type l (voter 1)

Red! (voter 2)

ok (voter 1)

Conversation 2 (Transparency):

May I guess, you will also again say you are of type h? (voter 2)

Cannot be the case that both are always h. (voter 2)

You can really give up on the chat. (voter 1)

OA3 Instructions to the experiment and screenshots

Instructions and screenshots for the transparency treatment (translation; original in German). Instructions and screens for the other treatments where very similar and are therefore omitted here. The original instructions can be obtained from the authors upon request.

Overview

Welcome to this experiment. We kindly ask you not to communicate with other participants during the experiment and to switch off your phones and other mobile devices.

At the end of the experiment you will be paid out in cash for your participation in today's session. The amount of your pay-off depends in parts on your decisions, on the decisions of other participants and on chance. For this reason it is important that you read the instructions carefully and understand them before the start of the experiment.

In this experiment all interactions between participants take place via the computers that you are sitting in front of. You will interact anonymously and your decisions will only be stored together with your random ID number. Neither your name nor names of other participants will be made public, not today and not in future written analyses.

Today's session consists of several rounds. **At the end, 4 rounds will be randomly selected and paid out. The rounds that are not chosen will not be paid out.** Your pay-off results from the points that you earn in the selected rounds, converted to Swiss Francs, plus your show-up fee of CHF 10. The conversion of points to Swiss Francs happens as follows: Every point is worth 15 cents, which means that

20 points = CHF 3.00.

Every participant will be paid out in private at the payment counter, so that no other participant can see how much you have earned.

Experiment

This experiment consists of 20 procedurally identical rounds. In each round a group decision has to be made that can be correct or wrong.

Two members in each group of three make the group decision (henceforth we will call them the voters). There are well and less well informed voters and the task of the third group member is to observe the decision process of the other two members and then to indicate the probability with which he thinks that the other group members are well or less well informed (henceforth, we will call this member the observer).

The higher the evaluation of the observer with respect to the level of information of a voter is, the higher is the pay-off to that voter in the round. The more accurate the evaluation of the observer with respect to the level of information of the voters is, the higher is his or her pay-off in the round. In addition, the observer receives a pay-off for correct group decisions.

The Group

In the **first** round you will be assigned a meta-group of 9 members. In the beginning of **every** round you will be randomly assigned to a **new group** which consists of randomly selected members of your meta-group. Every group has three members: 2 voters and 1 observer.

Whether you will be assigned the role of a voter or an observer, is randomly determined each round. The voters receive, again randomly, the labels “voter 1” and “voter 2”.

All interactions in a round take place within your group of three.

The Voters

There are two types of voters, well informed (type G) and (less well) informed (type I) voters. Of which type the group members are, is again determined randomly. With probability $\frac{1}{4}$ (or 25%) a voter receives good information which means he is of type G; with probability $\frac{3}{4}$ (or 75%) he receives less good information which means he is of type I.

Because the assignment of types to the voters is independent of the assignment to other voters, there can be two voters of type G, two voters of type I, or one of each type in a group.

The voters learn their type on the first screen of a round but not the type of the other voter in their group. The observer learns that he is an observer on the first screen but not the types of the voters in his group.

Later, after observing the behavior of the voters, it will be the task of the observer to estimate the probabilities that voter 1 and voter 2 are of type G.

The Jar

There are two jars: one red jar and one blue jar. The red jar contains 11 red and 9 blue balls, the blue jar 11 blue and 9 red balls. Each round one jar will be randomly selected.

The task of the voters is to vote on the color of the jar. Each jar has an equal probability of being selected, that is, it will be selected with 50% probability.

The Ball

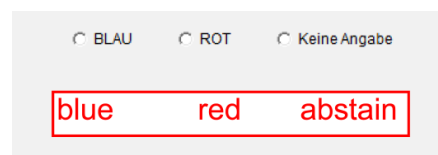
The well informed voters (type G) receive a ball with the actual color of the jar, that is they are directly informed about the color of the jar.

The informed voters (type I) receive a randomly drawn ball from the selected jar. They are not told the color of the jar. If there are two type I voters in a group, each of them receives a ball from the jar. Every ball in the jar has the same selection probability for the type I voters, that is for each voter of type I a ball is drawn from a jar containing 20 balls (11 with the color of the jar, 9 with the other color).

The voters learn the color of their ball on the first screen. Every voter only sees the color of his ball, not the color of the other voter's ball.

Communication

After learning their type and the color of their ball, the voters can communicate the color of their ball to the other voter in their group. They can also communicate the color that their ball did not have or stay silent. The communication is made through the following entry mask.



The image shows a user interface for communication. At the top, there are three radio buttons labeled "BLAU", "ROT", and "Keine Angabe". Below these, there is a red-bordered box containing the text "blue", "red", and "abstain" in red font, corresponding to the radio buttons above.

On the following screen the voters learn the message of the other voter in their group and have the option to chat with him. The chat happens via the following entry mask.

You can now chat with the other voter in your group. You will have the ID voter 1 and the other voter the ID voter 2 for this chat. Every entry has to be confirmed by pressing the enter button.

Sie können nun mit dem anderen Wähler in ihrer Gruppe chatten. Im Chat bekommen Sie die ID "Wähler 1" und der andere Wähler die ID "Wähler 2". Jede Texteingabe muss mit der Enter-Taste bestätigt werden.

Wähler 2: Hallo

You can enter arbitrary text messages into the blue entry field. Pay attention to confirm every entry by pressing the enter button to make it visible for the other voter. It will then appear in the grey field above.

The observer cannot participate in the communication but sees the messages of the two voters regarding the color of their ball as well as the chat.

Group Decision

After the communication stage the voters make their decision in a group vote.

So, if you are a voter, you have to vote either for blue or for red.

Once both voters have made their decision, the votes for blue and red are counted and the group decision results from the following rule:

- **If the color RED receives 2 votes, the group decision is RED**
- **If the color BLUE receives 1 or 2 votes, the group decision is BLUE**

That is for a group decision for blue only one vote is necessary while a group decision for red requires two votes.

Evaluation of the Observer

After the voters have cast their vote and the group decision is determined, the evaluator learns the group decision as well as the decisions of the individual voters in his group.

Moreover, he learns the true color of the jar, that is, whether the group decision and the individual decisions were correct or wrong.

On the same screen the observer can review the entire communication between the voters in his group once again.

If you are an observer, you now have to enter for each of the two voters the probability with which you believe that this voter is of type G.

To do so you enter a number between 0 and 100 which expresses your evaluation in percentage points. The entry mask looks as follows.

Geben Sie die Wahrscheinlichkeit in Prozent an, mit der Sie glauben, dass **Wähler 1** ein Typ G ist (eine Zahl zwischen 0 und 100):

Enter the probability in per cent, with which you believe that voter 1 is a high type.

The complete screen of the observer looks as follows (example screen).

Die Dose hat die Farbe **ROT** The color of the jar is red.

Die Wähler haben wie folgt abgestimmt The voters have voted as follows

Wähler 1 hat für folgende Farbe abgestimmt: **ROT** Voter 1 voted for RED. Wähler 2 hat für folgende Farbe abgestimmt: **ROT** Voter 2 voted for RED.

Die Gruppenentscheidung ist somit: **ROT** The resulting group decision is: RED.

Hier sehen Sie die Kommunikation zwischen den Wählern noch einmal Here you can once again review the communication between the voters.

Die Kugel, die Wähler 1 angab, ist: **ROT** Voter 1 indicated a red ball. Die Kugel, die Wähler 2 angab, ist: **ROT** Voter 2 indicated a red ball.

Wähler 2: Hallo

Geben Sie nun Ihre Einschätzung zu Wähler 1 und Wähler 2 an. Now, please, indicate your evaluation of voter 1 and voter 2.

Geben Sie die Wahrscheinlichkeit in Prozent an, mit der Sie glauben, dass **Wähler 1** ein Typ G ist (eine Zahl zwischen 0 und 100): Enter the probability in per cent, with which you believe that voter 1 is a high type.

Geben Sie die Wahrscheinlichkeit in Prozent an, mit der Sie glauben, dass **Wähler 2** ein Typ G ist (eine Zahl zwischen 0 und 100): Enter the probability in per cent, with which you believe that voter 2 is a high type.

Weiter

Pay-off in each Round

If you are a **voter** your pay-off is determined by the evaluation of the observer. If the observer believes that you are of type G with **X%** probability, you receive a pay-off of **2*X** points in

this round. This means that your pay-off directly depends on the probability with which the observer believes you are a **well-informed voter (type G)**.

If the observer has entered the probability 25%, for example, your pay-off is 50 points, if he has entered 50%, it is 100 points.

If you are an **observer** you receive a pay-off for correct group decisions and a pay-off for the accuracy of your evaluations of the types of the voters.

- If the group decision is RED and the jar is indeed RED, you as an observer receive **1 point**.
- If the group decision is BLUE and the jar is indeed BLUE, you as an observer receive **3 points**.
- If the group decision is wrong, you receive **0 points**, independently of the true color of the jar.

For your evaluation regarding the types of the voters you receive a pay-off between 0 and 100 points. It will be randomly determined whether you will be paid out for the evaluation of voter 1 or voter 2.

If you have evaluated both voters correctly with certainty (that is with 0 or 100%) (if you entered the probability 0 for both voters, for example, and both are indeed not of type G but of type I), you receive 100 points. If you are completely wrong (if both are of type G in the example) you receive 0 points.

The formula that determines your pay-off is a little complicated.

Put simple the formula assures **that it is best for you (gives you the highest expected pay-off) if you truthfully indicate the probability with which you believe that a voter is indeed of type G**. Every other evaluation lowers your expected pay-off.

If you believe, for example, that voter 1 in your group is of type G with 30% probability and voter 2 with 60% probability, it is best for you to enter exactly these values.

In case you want to know in more detail how your payoff is determined: for the evaluation of the randomly selected voter you receive:

$$100 - \frac{1}{100} (100 - \text{prob}(\text{voter is of type G}))^2, \text{ if this voter is of type G and}$$
$$100 - \frac{1}{100} (\text{prob}(\text{voter is of type G}))^2, \text{ if this voter is of type I,}$$

where $(\text{prob}(\text{voter is of type G}))$ is your indication of probability in percentage points that that voter is of type G. The resulting number is rounded up to a whole number and gives, together with your pay-off in case of a correct group decision, your pay-off in the round.

Remember: At the end of the experiment 4 rounds are randomly selected, the point incomes converted to Swiss Francs and paid out in private. The rounds that were not selected will not be paid out.

Questions?

Take your time to read the instructions carefully. If you have any questions, raise your hand. An experimental administrator will then come to your seat.

Screenshots *(not part of the instructions)*

First screen of a committee member

(The observer's first screen only informed the subject that he is an observer in that round.)

Sie sind ein Wähler vom
Typ I

Die Kugel, die Sie erhalten haben ist
ROT

You are a voter of type L

The ball that you have received is RED

You can now tell the second voter in your group what color your ball had. You can also stay silent or indicate the color that your ball did not have.

Sie können nun dem zweiten Wähler in ihrer Gruppe mitteilen, welche Farbe ihre Kugel hatte. Sie können auch schweigen, oder die Farbe, die ihre Kugel nicht hatte, als die Farbe ihrer Kugel angeben:

BLAU ROT Keine Angabe

blue red abstain

Weiter

Second screen of a committee member

Verbleibende Zeit [sec]: 72

You are a voter of type L

Sie sind ein Wähler vom
Typ I

The ball you received is red.

Die Kugel, die Sie erhalten haben ist
ROT

Die Kugel, die der andere Wähler in ihrer Gruppe angibt bekommen zu haben ist
ROT

The color of the ball of the other member of your group as indicated by him is red.

You can now chat with the other voter in your group. You will have the ID voter 1 and the other voter the ID voter 2 for this chat. Every entry has to be confirmed by pressing the enter button.

Sie können nun mit dem anderen Wähler in ihrer Gruppe chatten. Im Chat bekommen Sie die ID "Wähler 1" und der andere Wähler die ID "Wähler 2". Jede Texteingabe muss mit der Enter-Taste bestätigt werden.

Wähler 2: Hallo

Weiter

In the transparency treatment the principal could follow the chat in real time on a screen with a very similar layout. Under mild transparency and secrecy the principal just saw a waiting screen during communication.

Third screen of a committee member

You are a voter of type **L**
Sie sind ein Wähler vom
Typ I

The color of the ball of the other member of your group as indicated by him is red.
Die Kugel, die der andere Wähler in ihrer Gruppe angab bekommen zu haben ist
ROT

Die Kugel, die Sie erhalten haben ist
ROT

Here you can review the chat. Your ID was voter 1 and the other voter's ID was voter 2 for this chat.

Hier können Sie den Chat noch einmal nachlesen. Ihre ID war "Wähler 1" und die des anderen Wählers "Wähler 2".

Wähler 2: Hallo

Now make your final voting decision:
Treffen Sie nun Ihre Entscheidung für die Abstimmung:

BLAU ROT
blue red

Weiter

Evaluation screen of a principal in transparency treatment

Die Dose hat die Farbe **ROT** The color of the jar is red.

Die Wähler haben wie folgt abgestimmt The voters have voted as follows

Wähler 1 hat für folgende Farbe abgestimmt: **ROT** Voter 1 voted for RED. Wähler 2 hat für folgende Farbe abgestimmt: **ROT** Voter 2 voted for RED.

Die Gruppenentscheidung ist somit: **ROT** The resulting group decision is: RED.

Hier sehen Sie die Kommunikation zwischen den Wählern noch einmal
Here you can once again review the communication between the voters.

Die Kugel, die Wähler 1 angab, ist: **ROT** Voter 1 indicated a red ball. Die Kugel, die Wähler 2 angab, ist: **ROT** Voter 2 indicated a red ball.

Wähler 2: Hallo

Geben Sie nun Ihre Einschätzung zu Wähler 1 und Wähler 2 an.
Now, please, indicate your evaluation of voter 1 and voter 2.

Geben Sie die Wahrscheinlichkeit in Prozent an, mit der Sie glauben, dass **Wähler 1** ein Typ G ist (eine Zahl zwischen 0 und 100):

Geben Sie die Wahrscheinlichkeit in Prozent an, mit der Sie glauben, dass **Wähler 2** ein Typ G ist (eine Zahl zwischen 0 und 100):

Enter the probability in per cent, with which you believe that voter 1 is a high type.

Enter the probability in per cent, with which you believe that voter 2 is a high type.

Weiter

The evaluation screen had the same lay-out in the other two treatment but with several elements left out. Under mild transparency the communication part was left out. Under secrecy the individual votes and communication were left out and only one randomly selected committee member had to be evaluated.

Feedback screen for a voter at the end of a round

The screenshot shows a feedback screen for a voter at the end of a round. The screen is divided into several sections:

- Top section:** Two columns of information. The left column shows "Type of voter 1" (Typ von Wähler 1) and "Type L" (Typ I). The right column shows "Type of voter 2" (Typ von Wähler 2) and "Type H" (Typ G).
- Middle section:** A horizontal line separates the voter information from the group decision. Below the line, it states "The other voter in your group voted for: RED" (Der andere Wähler in ihrer Gruppe hat für folgende Farbe abgestimmt: Rot).
- Group decision:** It states "Die Gruppenentscheidung ist Rot" (The group decision is: RED).
- Jar color:** It states "Die Dose hat die Farbe Rot" (The color of the jar is: RED).
- Payoff:** It states "Ihre Rundenauszahlung (d.h. die Einschätzung des Beobachters verdoppelt) beträgt 60 Punkte" (Your payoff from this round (i.e. double the observer's evaluation) is 60 points).
- Bottom right:** A red button labeled "Weiter" (Next).

The feedback screens looked very similar for principals. At the end of the last round subjects saw a final screen which reported the rounds which were randomly selected to be paid out and the total earnings in Swiss Francs.